

AdaPT: An interactive procedure for multiple testing with side information

Lihua Lei* and William Fithian†

Department of Statistics, University of California, Berkeley

September 21, 2016

Abstract

We consider the problem of multiple hypothesis testing with generic side information: for each hypothesis H_i we observe both a p -value p_i and some predictor x_i encoding contextual information about the hypothesis. For large-scale problems, adaptively focusing power on the more promising hypotheses (those more likely to yield discoveries) can lead to much more powerful multiple testing procedures. We propose a general iterative framework for this problem, called the Adaptive p -value Thresholding (AdaPT) procedure, which adaptively estimates a Bayes-optimal p -value rejection threshold and controls the false discovery rate (FDR) in finite samples. At each iteration of the procedure, the analyst proposes a rejection threshold and observes partially censored p -values, estimates the false discovery proportion (FDP) below the threshold, and either stops to reject or proposes another threshold, until the estimated FDP is below α . Our procedure is adaptive in an unusually strong sense, permitting the analyst to use any statistical or machine learning method she chooses to estimate the optimal threshold, and to switch between different models at each iteration as information accrues.

1 Introduction

1.1 Interactive data analysis

In classical statistics we assume that the question to be answered, and the analysis to be used in answering the question, are both fixed in advance of collecting the data. Many modern applications, however, involve extremely complex data sets that may be collected without any specific hypothesis in mind. Indeed, very often the express goal is to explore the data in search of insights we may not have expected to find. A central challenge in modern statistics is to provide scientists with methods that are flexible enough to allow for exploration, but that nevertheless provide statistical guarantees for the conclusions that are eventually reported.

Selective inference methods blend exploratory and confirmatory analysis by allowing a search over the space of potentially-interesting questions, while still guaranteeing control of an appropriate Type I error rate such as a conditional error rate (e.g., Yekutieli, 2012; Lee et al., 2016; Fithian et al., 2014), familywise error rate (e.g., Tukey, 1994; Berk et al., 2013), or false discovery rate (e.g., Benjamini and Hochberg, 1995; Barber and Candès, 2015). However, most selective inference methods require that the selection algorithm be specified in advance, forcing a choice between either ignoring any difficult-to-formalize domain knowledge or sacrificing statistical validity guarantees.

*lihua.lei@berkeley.edu

†wfithian@berkeley.edu

Interactive data analysis methods relax the requirement of a pre-defined selection algorithm. Instead, they provide for an interactive analysis protocol between the analyst and the data, guaranteeing statistical validity as long as the protocol is followed. The two central questions in interactive data analysis are “what did the analyst know and when did she know it?” Previous methods for interactive data analysis involve randomization (Dwork et al., 2015; Tian and Taylor, 2015) to control the analyst’s access to the data at the time she decides what questions to ask.

This paper proposes an iterative, interactive method for multiple testing in the presence of *side information* about the hypotheses. We restrict the analyst’s knowledge by partially censoring all p -values smaller than a currently-proposed rejection threshold, and guarantee finite-sample FDR control by applying a version of the optional-stopping argument pioneered by Storey et al. (2004) and extended in Barber and Candès (2015); G’Sell et al. (2016); Li and Barber (2016a); Lei and Fithian (2016); Barber and Candès (2016).

1.2 Multiple testing with side information

In many areas of modern applied statistics, from genetics and neuroimaging to online advertising and finance, researchers routinely test thousands or millions of hypotheses at a time. For large-scale testing problems, perhaps the most celebrated multiple testing procedure of the modern era is the Benjamini–Hochberg (BH) procedure (Benjamini and Hochberg, 1995). Given n hypotheses and a p -value for each one, the BH procedure returns a list of rejections or “discoveries.” If R is the number of total rejections and V is the number of false rejections (rejections of true null hypotheses), the BH procedure controls the *false discovery rate* (FDR), defined as


$$\text{FDR} = \mathbb{E} \left[\frac{V}{\max\{R, 1\}} \right], \quad (1)$$

at a user-specified target level α . The random variable $V/\max\{R, 1\}$ is called the *false discovery proportion* (FDP).

The BH procedure is nearly optimal when the null hypotheses are exchangeable *a priori*, and nearly all true. In other settings, however, the power can be improved, sometimes dramatically, by applying prior knowledge or by learning from the data. For example, adaptive FDR-controlling procedures can gain in power by estimating the overall proportion of true nulls (Storey, 2002), applying priors to increase power using p -value weights (Benjamini and Hochberg, 1997; Genovese et al., 2006; Dobriban et al., 2015), grouping similar null hypotheses and estimating the true null proportion within each group (Hu et al., 2012), or exploiting a prior ordering to focus power on more “promising” hypotheses near the top of the ordering (Barber and Candès, 2015; G’Sell et al., 2016; Li and Barber, 2016a; Lei and Fithian, 2016).

In most large-scale testing problems, the null hypotheses do not comprise an undifferentiated list; rather, each hypothesis is associated with rich contextual information that could potentially help to inform our testing procedures. For example, Li and Barber (2016a) test for differential expression of 22,283 genes between a treatment and control condition for a breast cancer drug, with side information in the form of an ordering of genes from most to least “promising” using auxiliary data collected at larger dosages. Multiple testing procedures that exploit the ordering can reject hundreds of hypotheses while the BH procedure (which does not exploit the ordering) rejects none.

More generally, prior information could arise in more complex ways. For example, consider testing for association of 400,000 SNPs with each of 40 related diseases. If gene-regulatory relationships are known, then we might expect SNPs near related genes to be associated (or not) with related diseases, but without knowing ahead of time which gene-disease pairs are promising. Inspired by examples like this, Ignatiadis et al. (2016) and Li and Barber (2016b) have recently proposed a more general problem setting where, for each hypothesis H_i , $i \in 1, \dots, n$ we observe not only a p -value $p_i \in [0, 1]$ but also a predictor x_i



(a) $A_t = 4$ and $R_t = 11$ are the numbers of blue and red points respectively, leading to $\widehat{\text{FDP}}_t = (1 + 4)/11 \approx 0.45$. If $\widehat{\text{FDP}}_t \leq \alpha$, we stop and reject the red points; otherwise we choose a new threshold $s_{t+1}(x) \preceq s_t(x)$ and continue.

(b) Information available to the analyst when choosing $s_{t+1}(x)$ (A_t and R_t are also known). Each red and blue point is reflected across $p = 0.5$, leaving the analyst to impute which are the true p -values and which are the mirror images.

Figure 1: Illustration of one step of the AdaPT procedure with a univariate predictor.

lying in some generic space \mathcal{X} . The predictor is meant to capture some side information that might bear on H_i 's likelihood to be false, or on the power of p_i under the alternative, but the nature of this relationship is not fully known ahead of time and must be learned from the data.

In other situations, the “predictor” information could simply represent a measure of sample size or overall signal for testing the i th hypothesis, which could be informative about the power of the i th test to distinguish the alternative from the null. For example, if each p_i concerns a test for association between the i th SNP and a disease, then the overall prevalence of that SNP (in the combined treatment and control groups) can be used as prior information. Or, if p_i arises from a two-sample t -test, we could use the pooled variance of both groups as prior information (Ignatiadis et al., 2016).

1.3 AdaPT: a framework for FDR control

This paper presents a new framework for FDR control with generic side information, which we call *adaptive p-value thresholding* or AdaPT for short. Our method proceeds iteratively: at each step $t = 0, 1, \dots$, the analyst proposes a rejection threshold $s_t(x)$ and computes an estimator $\widehat{\text{FDP}}_t$ for the false discovery proportion for this threshold. If $\widehat{\text{FDP}}_t \leq \alpha$, she stops and rejects every H_i for which $p_i \leq s_t(x_i)$. Otherwise, she proposes a more stringent threshold $s_{t+1}(x) \preceq s_t(x)$ and moves on to the next iteration, where the notation $a(x) \preceq b(x)$ means $a(x) \leq b(x)$ for all $x \in \mathcal{X}$.

The estimator $\widehat{\text{FDP}}_t$ is computed by comparing the number R_t of rejections to the number A_t of p -values for which $p_i \geq 1 - s_t(x_i)$:

$$R_t = |\{i : p_i \leq s_t(x_i)\}|, \quad A_t = |\{i : p_i \geq 1 - s_t(x_i)\}|, \quad \text{and} \quad \widehat{\text{FDP}}_t = \frac{1 + A_t}{R_t \vee 1}.$$

Figure 1a illustrates the way $s_t(x)$ and $1 - s_t(x)$ partition the data into three regions; A_t is the number of points in the upper blue region and R_t is the number in the lower red region.

At each step t , the analyst can choose the next threshold $s_{t+1}(x)$ virtually however she chooses, with only two constraints. First, $s_{t+1}(x) \preceq s_t(x)$ as stated before. Second, the large and small p -values (the ones contributing to A_t and R_t) are partially masked. Specifically, at step t the analyst is allowed to observe A_t and R_t , as well as the entire sequence $(x_i, \tilde{p}_{t,i})_{i=1}^n$,

where

$$\tilde{p}_{t,i} = \begin{cases} p_i & s_t(x_i) < p < 1 - s_t(x_i) \\ \{p_i, 1 - p_i\} & \text{otherwise.} \end{cases} \quad (2)$$

Figure 1b illustrates what the analyst can see: each red and blue point from Figure 1a is shown along with its mirror image reflected across the midline $p = 0.5$.

We show in Section 3 that, roughly speaking, an ideal choice for $s_{t+1}(x)$ would be a level surface of the *local false discovery rate* (fdr), as a function of x and p :

$$\text{fdr}(p \mid x) = \mathbb{P}(H_i \text{ is null} \mid p_i = p, x_i = x).$$

Formally, $\text{fdr}(p \mid x)$ is unidentifiable from the data but, under reasonable assumptions, we can use a good proxy based on the conditional density $f(p \mid x)$.

In each step information is gradually revealed to the analyst as the threshold shrinks and more p -values are unmasked. Our procedure is adaptive in an unusually strong sense: provided that the two constraints are met, the analyst may apply any method she wants to select $s_{t+1}(x)$, consulting her own hunches or the intuition of domain experts, and can even switch between different methods as information accrues. Moreover, the analyst is under no obligation to describe, or even to fully understand, her update rule for choosing $s_{t+1}(x)$. In this sense, we say our method is fully *interactive* — the analyst’s behavior is arbitrary as long as she abides by a certain protocol for interacting with the algorithm.

The AdaPT procedure controls FDR at level α in finite samples provided that the null p -values are uniform and independent conditional on the non-null p -values. The proof relies on a pairwise exchangeability argument similar to the argument in Barber and Candès (2015).

Algorithm 1 summarizes the AdaPT procedure, using the generic sub-routine UPDATE to represent whatever process the analyst uses to select $s_{t+1}(x)$. Sections 3–4 discuss recommendations for a good UPDATE routine.

Algorithm 1 AdaPT

Input: predictors and p -values $(x_i, p_i)_{i=1,\dots,n}$, initialization s_0 , target FDR level α

Procedure:

```

1: for  $t = 0, 1, \dots$  do
2:    $\widehat{\text{FDP}}_t \leftarrow \frac{1+A_t}{R_t \vee 1}$ ;
3:   if  $\widehat{\text{FDP}}_t \leq \alpha$  then
4:     Reject  $\{H_i : p_i \leq s_t(x_i)\}$ ;
5:     Return  $s_t$ ;
6:   end if
7:    $s_{t+1} \leftarrow \text{UPDATE}((x_i, \tilde{p}_{t,i})_{i=1,\dots,n}, A_t, R_t, s_t)$ ;
8: end for
```

1.4 Related work

In recent work Ignatiadis et al. (2016) propose a different method for multiple testing with side information. They first bin the predictors into groups g_1, \dots, g_K , and then apply the weighted-BH procedure at level α with piecewise-constant weights; i.e., if $x_i \in g_k$, then $w_i = w(g_k)$. The weights $w(g_1), \dots, w(g_K)$ are chosen to maximize the number of rejections. This proposal is similar in spirit to the AdaPT procedure since it attempts to find optimal weights, but it is a bit more limited: first, binning the data may be difficult if the predictor space \mathcal{X} is multivariate or more complex; and second, their method is only guaranteed to control FDR asymptotically, as the number of bins stays fixed and the number of hypotheses in each bin grows to infinity. As a result, we must trust that n is large enough to support however many bins we have chosen to use. By contrast, AdaPT can use any machine-learning method to estimate $\hat{f}(p \mid x)$, and we can “overfit away” without

fear of compromising finite-sample FDR control (though overfitting can of course reduce our power). Another method is proposed by Du et al. (2014) when the covariate is an auxiliary univariate p-value derived by prior information. However, similar to Ignatiadis et al. (2016), it only controls FDR asymptotically under fairly strong conditions that the p-values are symmetrically distributed under the null and bounded by $\frac{1}{2}$ under the alternative.

Perhaps the procedure most closely related to ours is the *structure-adaptive BH algorithm* or SABHA (Li and Barber, 2016b). SABHA first censors the p -values below at a fixed level τ ($\tau = 0.5$ in their simulations), leading to censored p -values $p_i \mathbf{1}\{p_i > \tau\}$. Using these, they can estimate $\pi_1(x)$ as a function of x , then apply the weighted BH procedure of Genovese et al. (2006) with weights $\hat{\pi}_1(x_i)^{-1}$, at a corrected FDR level $\tilde{\alpha} = C\alpha$ (where $C < 1$ depends on the Rademacher complexity of the estimator $\hat{\pi}_1^{-1}$).

As the first procedure to provably control finite-sample FDR using generic feature information, SABHA represents a major step forward. However, AdaPT has several important advantages: First, even if $\hat{\pi}_1(x)$ estimates $\pi_1(x)$ consistently, the weights $\pi_1(x)^{-1}$ are not Bayes optimal as we show in Section 3; by contrast, our method estimates a Bayes optimal threshold. Second, the correction factor C makes the method conservative and restricts the available estimators $\hat{\pi}_1^{-1}$ to those with provably low Rademacher complexity. Third, AdaPT can use more information for learning: in later stages we will typically have $s_t(x_i) \ll 0.5$ and the masked p -values $\tilde{p}_{t,i}$ may be much more informative than $p_i \mathbf{1}\{p_i > 0.5\}$, especially when our goal is to estimate $f(p | x)$ for small values of p .

1.5 Outline

Section 2 defines the AdaPT procedure more formally and gives our main result: if the p -values are independent, AdaPT controls FDR at level α in finite samples. Section 3 explains why selection of $s_{t+1}(x)$ will typically operate by first estimating the conditional density $f(p | x)$ as a function of x , and Section 4 gives practical suggestions for update rules. Section 5 illustrates the AdaPT procedure’s power on the GEOQuery dosage response data, and Section 6 concludes.

2 The AdaPT procedure

2.1 Notation and assumptions

For each hypothesis H_i , $i = 1, \dots, n$ we observe $x_i \in \mathcal{X}$ and $p_i \in [0, 1]$. Let \mathcal{H}_0 denote the set of true null hypotheses. We will assume throughout that $(p_i)_{i \in \mathcal{H}_0}$ are mutually independent, and independent of $(p_i)_{i \notin \mathcal{H}_0}$ (see Section refsec:discussion for a discussion of how we might relax the independence assumption). Finally, for each $i \in \mathcal{H}_0$, we assume that p_i is either uniform or conservative in a sense we will define shortly.

Let \mathcal{F}_t represent the σ -field generated by all information available to the user at step t :

$$\mathcal{F}_t = \sigma((x_i, \tilde{p}_{t,i})_{i=1}^n, A_t, R_t).$$

The p -value masking ensures that $s_{t+1} \in \mathcal{F}_t^1$; together with the constraint $s_{t+1}(x) \preceq s_t(x)$ this ensures that $(\mathcal{F}_t)_{t=0,1,\dots}$ is a filtration; i.e., the information in \mathcal{F}_t only grows from t to $t + 1$.

To avoid trivialities we assume that, with probability 1, the analyst always reveals the next censored p -value within finitely many steps of the algorithm (indeed, there is no reason ever to update $s_{t+1}(x)$ in a way that reveals no new p -values). Thus, the stopping time \hat{t} is almost surely finite.

In many common settings, null p -values are conservative but not necessarily exactly uniform. For example, p -values from permutation tests are discrete, and p -values for composite

¹For simplicity we have implicitly ruled out the possibility that the analyst uses a randomized rule to update the threshold, but this restriction could be easily removed.

null hypotheses are often conservative if the true value of the parameter lies in the interior of the null.

Our method does not require uniformity, but the standard definition of conservatism — that $\mathbb{P}_{H_i}(p_i \leq a) \leq a$ for all $0 \leq a \leq 1$ — is *not* enough to guarantee FDR control. Instead, we say that a p -value p_i is *mirror-conservative* if

$$\mathbb{P}_{H_i}(p_i \in [a_1, a_2]) \leq \mathbb{P}_{H_i}(p_i \in [1 - a_2, 1 - a_1]), \quad \text{for all } 0 \leq a_1 \leq a_2 \leq 0.5. \quad (3)$$

If p_i is discrete, (3) means $p_i = 1 - a$ is at least as likely as $p_i = a$ for $a \leq 0.5$; if p_i has a continuous density, it means the density is at least as large at $1 - a$ as at a . Mirror-conservatism is not a consequence of conservatism (take $p_i = 0.1 + 0.9B$ where $B \sim \text{Bernoulli}(0.9)$), and neither does it imply conservatism (take $p_i = B$).

Permutation p -values are mirror-conservative, as are p -values for one-sided tests of univariate parameters with monotone likelihood ratio (with discrete p -values randomized to be uniform at the boundary between the null and alternative).

2.2 FDR control

We are now prepared to prove our main result: the AdaPT procedure controls FDR in finite samples. The proof relies on a similar optional stopping argument as the one presented in Lei and Fithian (2016) and Barber and Candès (2016) (themselves modifications of arguments in Storey et al. (2004) and Barber and Candès (2015)). Let V_t and U_t denote the numbers of null $p_i \leq s_t(x_i)$ and null $p_i \geq 1 - s_t(x_i)$, respectively. If the null p -values are uniform then, no matter how we choose $s_t(x)$ at each step, we will always have $V_t \approx U_t$ and $\widehat{\text{FDP}}_t > \frac{U_t}{R_t \vee 1} \approx \frac{V_t}{R_t \vee 1}$.

Lemma 1. *Suppose that, conditionally on the σ -field \mathcal{G}_{-1} , b_1, \dots, b_n are independent Bernoulli random variables with $\mathbb{P}(b_i = 1 \mid \mathcal{G}_{-1}) = \rho_i \geq \rho > 0$, almost surely. Also suppose that $\{1, \dots, n\} \supseteq \mathcal{C}_0 \supseteq \mathcal{C}_1 \supseteq \dots$, with each subset \mathcal{C}_{t+1} measurable with respect to*

$$\mathcal{G}_t = \sigma \left(\mathcal{G}_{-1}, \mathcal{C}_t, (b_i)_{i \notin \mathcal{C}_t}, \sum_{i \in \mathcal{C}_t} b_i \right)$$

If \hat{t} is an almost-surely finite stopping time with respect to the filtration $(\mathcal{G}_t)_{t \geq 0}$, then

$$\mathbb{E} \left[\frac{1 + |\mathcal{C}_{\hat{t}}|}{1 + \sum_{i \in \mathcal{C}_{\hat{t}}} b_i} \right] \leq \rho^{-1}.$$

Our Lemma 1 generalizes Lemma 1 in Barber and Candès (2016) and uses a very similar technical argument. The proof is given in the appendix. Using Lemma 1, we can give our main result:

Theorem 1. *Assume that the null p -values are independent of each other and of the non-null p -values, and the null p -values are uniform or mirror-conservative. Then the AdaPT procedure controls the FDR at level α .*

Proof. Let \hat{t} denote the step at which we stop and reject. Then

$$\text{FDP}_{\hat{t}} = \frac{V_{\hat{t}}}{R_{\hat{t}} \vee 1} = \frac{1 + U_{\hat{t}}}{R_{\hat{t}} \vee 1} \cdot \frac{V_{\hat{t}}}{1 + U_{\hat{t}}} \leq \alpha \frac{V_{\hat{t}}}{1 + U_{\hat{t}}},$$

where the last step follows from the stopping condition that $\widehat{\text{FDP}}_{\hat{t}} \leq \alpha$, and the fact that $U_t \leq A_t$. We will finish the proof by establishing that $\mathbb{E}[V_{\hat{t}}/(1 + U_{\hat{t}})] \leq 1$, using Lemma 1.

Let $m_i = \min\{p_i, 1 - p_i\}$ and $b_i = \mathbf{1}\{p_i \geq 0.5\}$, so $p_i = b_i(1 - m_i) + (1 - b_i)m_i$. Then knowing b_i and m_i is equivalent to knowing p_i . Let $\mathcal{C}_t = \{i \in \mathcal{H}_0 : p_i \notin (s_t(x_i), 1 - s_t(x_i))\}$, representing the null p -values that are *not* visible to the analyst at time t . Then,

$$U_t = \sum_{i \in \mathcal{C}_t} b_i, \quad \text{and} \quad V_t = \sum_{i \in \mathcal{C}_t} (1 - b_i) = |\mathcal{C}_t| - U_t.$$

Further, define the σ -fields

$$\mathcal{G}_{-1} = \sigma((x_i, m_i)_{i=1}^n, (b_i)_{i \notin \mathcal{H}_0}), \quad \text{and} \quad \mathcal{G}_t = \sigma(\mathcal{G}_{-1}, \mathcal{C}_t, (b_i)_{i \notin \mathcal{C}_t}, U_t).$$

The assumptions of independence and mirror-conservatism guarantee $\mathbb{P}(b_i = 1 \mid \mathcal{G}_{-1}) \geq 0.5$ almost surely for each $i \in \mathcal{H}_0$, with the b_i conditionally independent.

Next, note that $\mathcal{F}_t \subseteq \mathcal{G}_t$ because $p_i \in \mathcal{G}_t$ for each $p_i \notin (s_t(x_i), 1 - s_t(x_i))$, and

$$A_t = U_t + |\{i \notin \mathcal{H}_0 : p_i \geq 1 - s_t(x_i)\}|,$$

and $R_t \in \mathcal{G}_t$ by a similar argument. It follows that $\hat{t} = \min\{t : \widehat{\text{FDP}}_t \leq \alpha\}$ is a stopping time with respect to \mathcal{G}_t ; furthermore, $\mathcal{C}_{t+1} \in \mathcal{F}_t \subseteq \mathcal{G}_t$ by assumption.

As a result, we can apply Lemma 1 to obtain

$$\text{FDR} \leq \alpha \mathbb{E} \left[\frac{V_{\hat{t}}}{1 + U_{\hat{t}}} \right] = \alpha \mathbb{E} \left[\frac{1 + |\mathcal{C}_{\hat{t}}|}{1 + U_{\hat{t}}} - 1 \right] \leq \alpha(2 - 1) = \alpha.$$

□

The main technical point of departure for our method is that the optional stopping argument is not merely a technical device to prove FDR control for a fixed algorithm like the BH, Storey-BH, or Knockoff+ procedures. Instead, we push the optional-stopping argument to its limit, allowing the analyst to interact with the data in a much more flexible and adaptive way. Sections 6.1–6.2 further investigate the connection to knockoffs.

3 Admissible thresholding rules

Although the AdaPT procedure controls FDR no matter how we update the threshold, its power depends on the quality of the updates. This section concerns the question of what thresholds we would choose if we had perfect knowledge of the data-generating distribution, with Section 4 discussing suggestions for learning optimal thresholds from the data. As we will see, under mild conditions, the Bayes-optimal rejection thresholds are the level surfaces of the *local false discovery rate* (fdr), defined as the probability that a hypothesis is null conditional on its p -value. The local FDR was first discussed by Efron et al. (2001); see also Efron (2007). A similar result is obtained by Storey (2007) under a different framework.

3.1 The two-groups model and local false discovery rate

To begin, we assume a *two-groups model* conditional on the predictors x_i . Letting $H_i = 0$ if the i th null is true and $H_i = 1$ otherwise, we assume:

$$\begin{aligned} H_i \mid x_i &\sim \text{Bernoulli}(\pi_1(x_i)) \\ p_i \mid H_i, x_i &\sim \begin{cases} f_0(p \mid x_i) & \text{if } H_i = 0 \\ f_1(p \mid x_i) & \text{if } H_i = 1 \end{cases}. \end{aligned}$$

In addition, we assume that (x_i, H_i, p_i) are independent for $i = 1, \dots, n$. Unless otherwise stated we will assume for simplicity that both f_0 and f_1 are continuous densities, with $f_0(p \mid x) \equiv 1$ (null p -values are uniform) and $f_1(p \mid x)$ non-increasing in p (smaller p -values imply stronger evidence against the null). Furthermore, define the conditional mixture density

$$f(p \mid x) = (1 - \pi_1(x)) f_0(p \mid x) + \pi_1(x) f_1(p \mid x) = 1 - \pi_1(x) + \pi_1(x) f_1(p \mid x),$$

and the conditional local false discovery rate

$$\text{fdr}(p \mid x) = \mathbb{P}(H_i \text{ is null} \mid x_i = x, p_i = p) = \frac{1 - \pi_1(x)}{f(p \mid x)}.$$

Note that we never observe H_i directly. Thus, while f is identifiable from the data, π_1 and f_1 are not: for example, $\pi_1 = 0.5, f_1(p | x) = 2p$ and $\pi_1 = 1, f_1(p | x) = p + 0.5$ result in exactly the same mixture density. Unless $f_1(p | x)$ is known *a priori*, we can make the conservative identifying assumption that

$$1 - \pi_1(x) = \inf_{p \in [0,1]} f(p | x) = f(1 | x),$$

attributing as many observations as possible to the null hypothesis. This approximation is very good when $\text{fdr}(1 | x) \approx 1$, which is reasonable in many settings. Thus, any estimate \hat{f} of the mixture density translates to a conservative estimate $\widehat{\text{fdr}}(p | x) = \hat{f}(1 | x) / \hat{f}(p | x)$.

3.2 Admissible thresholds

Let ν be a probability measure on \mathcal{X} and define a random variable $X \sim \nu$. Similar to Sun et al. (2015), for any thresholding rule $s(x)$, we define the global FDR as

$$\text{FDR}(s; \nu) = \mathbb{P}(H = 0 | H \text{ is rejected}) = \mathbb{P}(H = 0 | P \leq s(X))$$

where H and P are a hypothesis and p -value distributed according to the two-groups model. The power is defined in a similar fashion as

$$\text{Pow}(s; \nu) = \mathbb{P}(H \text{ is rejected} | H = 1) = \mathbb{P}(P \leq s(X) | H = 1).$$

Sun et al. (2015) formulates a compound decision-theoretic framework by defining a Bayesian-type loss function. Instead, we propose a Neyman-Pearson type framework, i.e.

$$\max_s \text{Pow}(s; \nu) \quad \text{s.t.} \quad \text{FDR}(s; \nu) \leq \alpha. \quad (4)$$

We call a thresholding rule s *admissible* if s solves the above optimization problem for some $\alpha \in [0, 1]$.

Next, define

$$\begin{aligned} Q_0(s) &= \mathbb{P}(P \leq s(X), H = 0) = \int_{\mathcal{X}} F_0(s(x)|x)(1 - \pi_1(x))\nu(dx) \\ Q_1(s) &= \mathbb{P}(P \leq s(X), H = 1) = \int_{\mathcal{X}} F_1(s(x)|x)\pi_1(x)\nu(dx), \end{aligned}$$

where F_0 and F_1 are the cumulative distribution functions under the null and alternative. We can simplify (4) as

$$\max_s \frac{Q_1(s)}{\mathbb{P}(H = 1)} \quad \text{s.t.} \quad \frac{Q_0(s)}{Q_0(s) + Q_1(s)} \leq \alpha \quad (5)$$

$$\iff \min_s -Q_1(s) \quad \text{s.t.} \quad -\alpha Q_1(s) + (1 - \alpha)Q_0(s) \leq 0 \quad (6)$$

$$\begin{aligned} \iff \min_s \int_{\mathcal{X}} -F_1(s(x)|x)\pi_1(x)\nu(dx) \\ \text{s.t.} \quad \int_{\mathcal{X}} \left\{ -\alpha F_1(s(x)|x)\pi_1(x) + (1 - \alpha)F_0(s(x)|x)(1 - \pi_1(x)) \right\} \nu(dx) \leq 0. \end{aligned} \quad (7)$$

The corresponding Lagrangian function can be written as

$$L(s; \lambda) = \int_{\mathcal{X}} \left\{ -(1 + \lambda\alpha)F_1(s(x)|x)\pi_1(x) + \lambda(1 - \alpha)F_0(s(x)|x)(1 - \pi_1(x)) \right\} \nu(dx). \quad (8)$$

Let s^* be the optimum, then KKT condition (under regularity conditions) implies that

$$\begin{aligned} (1 + \lambda\alpha)f_1(s^*(x)|x)\pi_1(x) &= \lambda(1 - \alpha)f_0(s^*(x)|x)(1 - \pi_1(x)) \\ \implies \text{fdr}(s^*(x)|x) &= \frac{1 + \lambda\alpha}{1 + \lambda}. \end{aligned} \quad (9)$$

In other words, all admissible thresholding rules are level surfaces of local FDR. Theorem 2 formalizes the above derivation by clarifying the regularity conditions.

Theorem 2. *Assume that*

- (i) $f_1(p | x_i)$ is continuously non-increasing and $f_0(p | x_i)$ is continuously non-decreasing and uniformly bounded away from ∞ ;
- (ii) ν is a discrete measure supported on $\{x_1, \dots, x_n\}$ with $\nu(\{x_i : \widehat{\text{fdr}}(0 | x_i) < \alpha, f(0 | x_i) > 0\}) > 0$.

Then (4) has at least a solution, and all solutions are level surfaces of $\widehat{\text{fdr}}(p | x)$.

In practice, any conservative null distribution (stochastically dominated by $U([0, 1])$) with positive density at zero satisfies condition (i). The monotonicity of f_1 is also valid since smaller p-values imply stronger evidence against null. In condition (ii), the assumption on the support is reasonable since we treat $\{x_i : i = 1, \dots, n\}$ as fixed and hence only the quantities associated with these values are of interest. We believe it can be relaxed to more general measures and will not discuss it due to the technical complication. On contrast, the second requirement is necessary since it implies the feasibility of the problem. If the local FDR is above α almost everywhere, no thresholding rule is able to control FDR at α . As mentioned above, we can set s as the level surfaces of $\widehat{\text{fdr}}(p | x) = \hat{f}(1 | x)/\hat{f}(p | x)$ given some estimator $\hat{f}(p | x)$. The next section discusses estimation of $\hat{f}(p | x)$.

4 Implementation

Having shown that level surfaces of the local FDR are optimal, we now turn to estimation of $\widehat{\text{fdr}}(p | x)$. This section discusses a flexible framework for conditional density estimation that can perform favorably when no domain-specific expertise can be brought to bear.

More generally, we should model the data using as much domain-specific expertise as possible. We emphasize once more that, no matter how misspecified our model is, no matter how misguided our priors are (if we use a Bayesian method), no matter how we select a model or tuning parameter, or how much that selection biases our resulting estimate of local FDR, the AdaPT procedure nevertheless controls global FDR. Thus, there is every reason to be relatively aggressive in choosing a modeling strategy.

4.1 Conditional density estimation and generalized linear modeling

Generically, we can convert our conditional density estimation problem into a generalized linear modeling problem using Lehmann alternatives, where we assume $F_0(p | x) = p$ and $F_1(p | x) = p^{1/\mu(x)}$ for $\mu(x) \in (0, 1]$. Transforming the p -values to $y_i = -\log p_i$ (note $y_i \sim \text{Exp}(1)$ if $p_i \sim U[0, 1]$), and letting $\phi(x) \in \mathbb{R}^d$ denote some featurization, this leads to the two-group Gamma GLM mixture model for y_i :

$$H_i | x_i \sim \text{Bernoulli}(\pi_1(x_i)), \quad \text{with } \log \frac{\pi_1(x)}{1 - \pi_1(x)} = \theta' \phi(x), \quad \text{and} \quad (10)$$

$$y_i | x_i, H_i \sim \begin{cases} \text{Exp}(\mu(x_i)) & \text{if } H_i = 1 \\ \text{Exp}(1) & \text{if } H_i = 0 \end{cases}, \quad \text{with } \mu(x)^{-1} = \beta' \phi(x).$$

Note that larger values of μ indicate regions with more signal. At step t , the y_i values are partially masked (owing to the partial masking of p_i) and all of the H_i values are unknown, but we can efficiently estimate the coefficients of (10) by using the EM algorithm, which imputes y_i and H_i at each E-step and estimates a logistic regression and weighted Gamma GLM at each M-step. Details are given in Appendix A. Note it is not necessary to use the same featurization for predicting both π_1 and μ .

The Gamma GLM model (10) provides the starting point for an extremely flexible and extensible modeling framework. More generally, we could estimate a GLM with a lasso or ridge penalty, generalized additive model, random forest (with Gamma regression trees as

the base learner), hierarchical Bayesian models, collaborative filtering methods, and many more.

4.2 Updating the threshold

Theorem 2 suggests that our updated threshold s_{t+1} should approximate a level surface of $\text{fdr}(p \mid x)$. For the model (10), level surfaces of the local FDR are given by²

$$c = \frac{f(1|x)}{f(s(x)|x)} = \frac{\pi_1(x) \cdot \frac{1}{\mu(x)} + 1 - \pi_1(x)}{\pi_1(x) \cdot \frac{1}{\mu(x)} \cdot s(x)^{\frac{1}{\mu(x)} - 1} + 1 - \pi_1(x)}$$

$$\implies s(x) = \left(\frac{1}{c} + \mu(x) \cdot \frac{1 - \pi_1(x)}{\pi_1(x)} \cdot \frac{1 - c}{c} \right)^{\frac{\mu(x)}{1 - \mu(x)}}.$$

Denote $\tilde{s}(x; c)$ by the above surface and for any chosen c , we can evolve s_t by

$$s_{t+1}(x) = \min\{s_t(x), \tilde{s}(x; c)\}, \quad (11)$$

where the minimum is taken to meet the requirement that $s_{t+1}(x) \leq s_t(x)$. Note that a higher level surface (larger c) will typically give a higher $\widehat{\text{FDP}}_t$ and vice versa. Unless computational efficiency is at a premium, it is better to force the procedure to be patient since more information can be gained after each update and the learning step can be more accurate. In other words, we shall choose a large c such that $s_{t+1}(x)$ only deviates from $s_t(x)$ slightly. In this article we propose a simple procedure to achieve this. It chooses c as

$$\hat{c} = \min \left\{ c \in [0, 1] : \frac{\sum_i \min\{s_t(x_i), \tilde{s}(x_i; c)\}}{\sum_i s_t(x_i)} \geq 1 - \delta \right\} \quad (12)$$

where δ is a patience parameter and we set it to be 0.05 as a default. More details of choosing δ is given in Appendix C. \hat{c} could be solved by binary search efficiently.

On the other hand, it could happen at early stages, when the revealed information is limited, that the estimates of local FDR is inaccurate. If (11) is used directly, although patient globally it is still intolerant at those points at which the level surfaces are mistakenly low due to the lack of accuracy. To avoid this, we propose the following modification which guarantees that $s_t(x)$ will never shrink too fast:

$$s_{t+1}(x) = \max\{(1 - \delta)s_t(x), \min\{s_t(x), \tilde{s}(x; \hat{c})\}\} \quad (13)$$

where δ is the same as above and \hat{c} is determined by (12). Finally, to avoid the unexpected case that $s_{t+1}(x) = s_t(x)$, which could happen when δ is too small, we create a buffer storing the s -curve for the last B periods, namely $\{s_{t-B}(x), \dots, s_t(x)\}$. If all of them all almost equal then we force $s_{t+1}(x)$ to be reduced as $s_{t+1}(x) = (1 - \delta)s_t(x)$. We set $B = 5$ as a default.

4.3 Other Issues

First we discuss the initialization of AdaPT. As shown in Algorithm 1, AdaPT starts from some curve $s_0(x)$ and then slowly update it. If the hypotheses are not ordered, then we can simply set $s_0(x) \equiv s_0$ with $s_{0,1} < 0.5$. A larger $s_{0,1}$ is conceptually preferred since the procedure is more patient. We found that $s_0 = 0.45$ is a consistently good choice.

If the hypotheses is ordered and believed to be arranged from most promising to least promising, we recommend using a piecewise linear s_0 with

$$s_0(x) = s_{0,1}I(x \leq \gamma n) + s_{0,2}I(x > \gamma n).$$

²Because estimating the model (10) does yield estimates of f_1 and π_1 , we could use these directly to estimate fdr , but we find that estimates using $\hat{f}(1 \mid x)/\hat{f}(p \mid x)$ are more stable.

where $0.5 > s_{0,1} > s_{0,2} > 0$ and $\gamma \in (0, 1]$. In other words, we treat the first γ -fraction of hypotheses as promising ones and start from a high threshold and the last $(1 - \gamma)$ of hypotheses as noise and start from a low threshold. The performance of this initialization is not quite different from the above one for small α but could save computational cost significantly. We found that $s_{0,1} = 0.45, s_{0,2} = 0.05, \gamma = 0.5$ is a good choice consistently.

Another issue is computation. In applications, researchers might be interested in the discoveries with a range of α 's. Since AdaPT is a computational-intensive procedure, it is costly to run the procedure for each single α . Fortunately, AdaPT is able to produce the rejections for all $\alpha \in [0, \hat{q}_0]$ with only one pass where \hat{q}_0 is the estimate of local FDR with thresholding rule s_0 . By setting δ (in (12) and (13)) sufficiently small, \hat{q}_t will evolve slowly. For each $\alpha \in [0, \hat{q}_0]$, let $t(\alpha) = \min\{t : \hat{q}_t \leq \alpha\}$, then the thresholding rule $s_{t(\alpha)}$ controls FDR at level α . In addition, a large tolerance parameter δ makes the procedure less patient and hence more computationally efficient. This motivates us to apply a large δ when $\widehat{\text{fdp}} \gg \alpha$ and reduce it after $\widehat{\text{fdp}}$ is closed to α .

5 GEOQuery data

To illustrate the power of the AdaPT procedure, we apply it to the GEOQuery data (Davis and Meltzer, 2007), which has been analyzed repeatedly as a benchmark for ordered testing procedures Li and Barber (2016a); Lei and Fithian (2016); Li and Barber (2016b). We use Algorithm 2 in Appendix A. This dataset consists of gene expression measurements for $n = 22283$ genes, in response to estrogen treatments in breast cancer cells for five groups of patients, with different dosage levels and 5 trials in each. The task is to identify the genes responding to a low dosage. The p -values p_i for gene i is obtained by a one-sided permutation test which evaluates evidence for a change in gene expression level between the control group (placebo) and the low-dose group. $\{p_i : i = 1, \dots, n\}$ are then ordered according to permutation t -statistics comparing the control and low-dose data, pooled, against data from a higher dosage (with genes that appear to have a strong response at higher dosages placed earlier in the list).

We consider three orderings: first, a stronger (more informative) ordering based on a comparison to the highest dosage; second, a weaker (less informative) ordering based on a comparison to a medium dosage; and third, the original ordering of the data using no auxiliary data. Let $\sigma_S(i)$ and $\sigma_W(i)$ denote respectively the permutations of $i = 1, \dots, n$ given by the stronger and weaker orderings. Further details on these two orderings can be found in Li and Barber (2016a) and Li and Barber (2016b). We write the p -values, thus reordered, as $p_i^S = p_{\sigma_S(i)}$ and $p_i^W = p_{\sigma_W(i)}$. Once the data are re-ordered, we can apply either a method that ignores the ordering altogether, or an ordered testing procedure, or a testing procedure that uses generic side information, using the index of the re-ordered p -values as a univariate predictor.

We compare AdaPT against seven other methods³: BH procedure (Benjamini and Hochberg, 1995), Storey's BH procedure with threshold $\lambda = 0.5$ (Storey et al., 2004), SABHA with the algorithm described in section 4.2 of Li and Barber (2016b), SeqStep with parameter $C = 2$ (Barber and Candès, 2015), ForwardStop (G'Sell et al., 2016), the accumulation test with the HingeExp function and parameter $C = 2$ (Li and Barber, 2016a), Adaptive SeqStep with $s = q$ and $\lambda = 1 - q$ (Lei and Fithian, 2016) and Independent Hypothesis Weighting with number of bins and folds set as default (Ignatiadis et al., 2016). Figure 2 shows the number of discoveries with different target FDR levels. We only show the range of α from 0.05 to 0.3 since it is rare to allow FDR to be above 0.3 in practice. The implementation details of AdaPT is given in Appendix A. $\phi(x)$ is taken to be the basis functions of natural spline with 8 equi-spaced knots.

³The code is available at <https://github.com/lihuallei71/AdaPT/>




Figure 2: Number of discoveries by each method at a range of target FDR levels α from 0.01 to 0.30. Each panel plots the results for an ordering, ranging from no ordering to highly informative.

The left panel of Figure 2 corresponds to the original ordering, and shows that AdaPT is the only method with power to detect signals when $\alpha < 0.3$, yielding 722 discoveries when $\alpha = 0.29$. In addition, AdaPT has 146 discoveries when α is as small as 0.2. Figure 3 shows the $s(x)$ at $\alpha = 0.2$ and $\alpha = 0.3$ in the upper panels. The bottom panels plot the contours of estimated local FDR under model (10). As discussed in section 3, the thresholding rule $s(x)$ should exhibit similar form as local FDR. This is justified in both panels of Figure 2. Although no auxiliary data is used for ordering, we can still observe a bump at the beginning of the list, which contributes a large number of discoveries. This is partially because that AdaPT is able to detect spatial patterns and uses it to improve power.

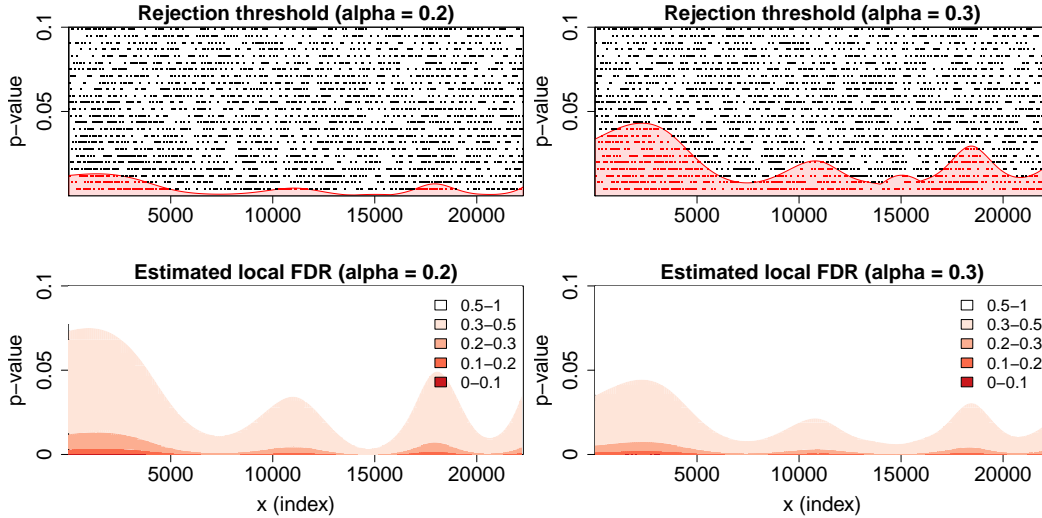


Figure 3: Results for GEOquery data with original ordering of p-values with $\alpha = 0.2$ (left) and $\alpha = 0.3$ (right): (top) the dots represent the p-values and the red dots are rejected ones. The red curve is the thresholding rule $s(x)$; (bottom) the contour plots of estimated local FDR.

The next two panels of Figure 2 shows that AdaPT consistently more powerful than all other methods, especially when α is small. In many applications, the target FDR level should be as low as 0.05. For moderately informative ordering, AdaPT and Adaptive SeqStep are the only methods which have non-zero discoveries at $\alpha = 0.05$ and AdaPT has 299 discoveries, which is about twice as many discoveries as Adaptive SeqStep. For highly infomative ordering, AdaPT has 874 discoveries, while Adaptive SeqStep has 480 ones and

Accumulation test with HingeExp function has 349 ones. Figure 4 and Figure 5 plot the thresholding rules and estimated signal strength for p-values with moderately informative ordering and p-values with highly informative ordering, respectively. It can be seen from the bottom panels that the evidence to be non-null has an obvious decreasing trend when the ordering is used. Moreover, the highly informative ordering indeed sorts the p-values better than the moderately informative ordering. For the former, the thresholding rule is fairly monotone while it has a small bump at $i \approx 5000$ for the latter. In both cases, most discoveries are from the first 5000 genes in the list.

In summary, these plots show a strong data adaptivity of AdaPT, which can also learn the local structure of data while controlling FDR. Moreover, it provides a quantitative way, by estimated signal strength, to evaluate the quality of ordering, which is the major concern in ordered testing problems (Li and Barber, 2016a; Lei and Fithian, 2016; Li and Barber, 2016b).




Figure 4: Results for GEOquery data with moderate ordering of p-values, i.e. $\{p_{(i)}^{(1)}\}$, with $\alpha = 0.05$ (left) and $\alpha = 0.1$ (right): (top) the dots represent the p-values and the red dots are rejected ones. The red curve is the thresholding rule $s(x)$; (bottom) the contour plots of estimated local FDR.

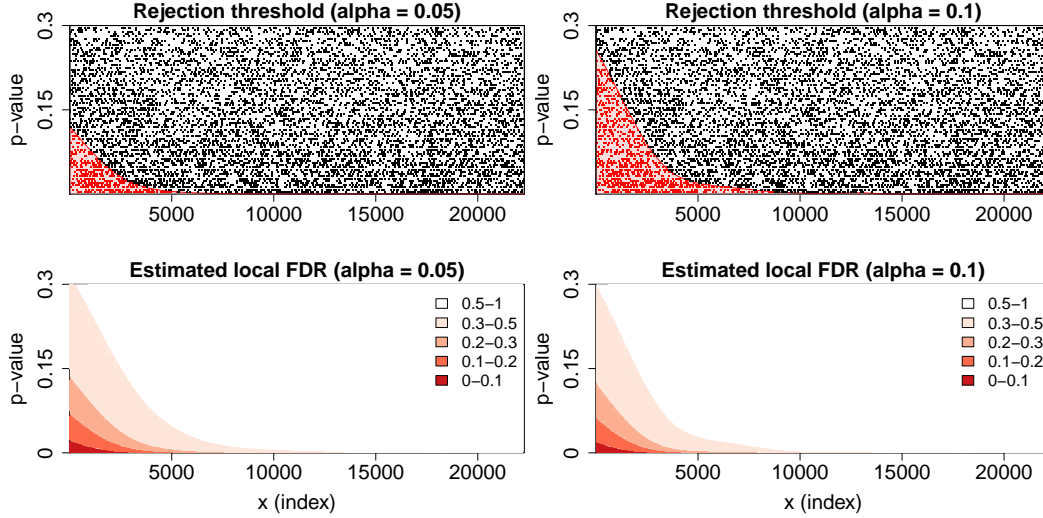


Figure 5: Results for GEOQuery data with moderate ordering of p -values, i.e. $\{p_{(i)}^{(2)}\}$, with $\alpha = 0.05$ (left) and $\alpha = 0.1$ (right): (top) the dots represent the p -values and the red dots are rejected ones. The red curve is the thresholding rule $s(x)$; (bottom) the contour plots of estimated local FDR.

6 Discussion

We have proposed the AdaPT procedure, a general iterative framework for multiple testing with side information. Using partially masked p -values, we estimate a family of optimal and increasingly stringent rejection thresholds, which are level surfaces of the local FDR. We then monitor an estimator of FDP to decide which threshold to use, updating our estimates as we unmask more p -values and gain more information.

Our method is interactive in that it allows the analyst to use an arbitrary method for estimating the local FDR, and to consult her intuition to change models at any iteration, even after observing most of the data. No matter what the analyst does or how badly she overfits the data, FDR is still controlled at the advertised level (though power could be adversely affected by overfitting). We show using the GEOQuery data that AdaPT can give significant power improvements over current state-of-the-art methods.

6.1 Extension to dependent data using knockoffs

It would also be interesting to attempt to relax our restriction that the p -values must be independent. In the absence of some modification, our AdaPT procedure does not control FDR in finite samples for dependent p -values. In particular, there is a danger of “overfitting” to local random effects that shared by nearby hypotheses: to the AdaPT procedure, such random effects are treated as signal to discover.

It could be interesting to pursue a hybrid method using ideas from AdaPT and Knockoff+ procedures in the case where the p -values arise from multivariate Gaussian test statistics. For example, in the two-tailed case we could have

$$z_i \sim \mathcal{N}(\mu_i, \sigma^2), \quad p_i = 2 \min\{\Phi(z_i), 1 - \Phi(z_i)\}, \quad \text{for } i = 1, \dots, n,$$

but with nonzero correlations between the z_i values.

Then we could first compute randomized “knockoff” z -values $z_i^* \sim \mathcal{N}(0, \sigma^2)$, using the methods of Barber and Candès (2015) to ensure that the distribution of (z, z^*) is invariant to permutations that swap z_i with z_i^* for any subset of null indices $i \in \mathcal{H}_0$. The same

pairwise exchangeability property would then apply to the induced “knockoff” p -values $p_i^* = 2 \min\{\Phi(z_i^*), 1 - \Phi(z_i^*)\}$.

Then, we could mask the p -values via

$$\tilde{p}_{t,i} = \begin{cases} p_i & \min\{p_i, p_i^*\} > s_t(x) \\ \{p_i, p_i^*\} & \text{otherwise} \end{cases},$$

and compute an FDP estimate based on the masked p -values via

$$R_t = |\{i : p_i \leq p_i^* \wedge s_t(x_i)\}|, \quad A_t = |\{i : p_i^* \leq p_i \wedge s_t(x_i)\}|, \quad \text{and} \quad \widehat{\text{FDP}}_t = \frac{1 + A_t}{R_t \vee 1}.$$

Note that the knockoff construction would typically guarantee $\mathbb{P}(\text{any } p_i = p_i^*) = 0$. Using a similar pairwise exchangeability argument it should be possible to show that this hybrid method will control FDR in finite samples.

6.2 Connection to knockoffs in the independence case

By focusing on the independence case, the above construction further illuminates the relationship between AdaPT and the Knockoff+ procedure. Suppose that the z_i represent estimates for regression coefficients in a regression with an orthonormal design matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ and response $y \sim \mathcal{N}_n(\mathbf{X}\beta, I_n)$, for $n \geq 2d$. Then we could construct the knockoff z -values z_i^* by introducing d more artificial variables, giving

$$(z, z^*) = [\mathbf{X} \ \mathbf{X}^*]' y \sim \mathcal{N}_{2d} \left(\begin{pmatrix} \beta \\ 0 \end{pmatrix}, I_{2d} \right), \quad \text{with} \quad [\mathbf{X} \ \mathbf{X}^*]' [\mathbf{X} \ \mathbf{X}^*] = I_{2d}.$$

If, for example, we used the LASSO for model selection as proposed in Barber and Candès (2015), it can be shown that the Knockoff+ procedure corresponds exactly to the hybrid proposal of Section 6.1, but with $s_t(\cdot)$ a constant function at each step. More generally, prior information about the different variables could inform the model-selection procedure, about which Knockoff+ is agnostic.

Thus, after transforming (z, z^*) to (p, p^*) , the setup of knockoffs looks similar to the setup for the standard AdaPT procedure. The two most salient differences are:

1. AdaPT allows for iterative interaction between the analyst and data, allowing the analyst to update her local FDR estimates as information accrues. By contrast, Knockoff+ as described by Barber and Candès (2015) does not allow for such interaction (though it could, and this is a potentially interesting avenue for extending knockoffs).
2. Unlike Knockoff+, AdaPT introduces no extra randomness into the problem. This is because AdaPT uses pairwise exchangeability of p_i with the “mirror image” p -value $1 - p_i$ instead of the independent “knockoff” p -value $p_i^* \sim U[0, 1]$. Thus, as a statistical procedure AdaPT respects the sufficiency principle: for any choice of UPDATE subroutine, the AdaPT result is a deterministic function of the original data.

6.3 Extension: estimating local FDR

In addition to returning a list of rejections that is guaranteed to control the global FDR, most implementations of AdaPT will also return estimates, for each rejected hypothesis, of the local FDR,

$$\widehat{\text{fdr}}(p_i \mid x_i) = \widehat{\mathbb{P}}(H_i \text{ is null} \mid x_i, p_i).$$

If we have reasonably high confidence in the model we have used to produce these estimates, they may provide the best summary of evidence against the individual hypothesis H_i . By contrast, the significance level for global FDR only summarizes the strength of evidence against the entire list of rejections, taken as a whole. Indeed, it is possible to

construct pathological examples where $\text{fdr}(p_i | x_i) = 1$ for some of the rejected H_i , despite controlling FDR at some level $\alpha \ll 1$. Even apart from such perversities, it will typically be the case that $\widehat{\text{fdr}}(p_i | x_i) > \alpha$ for many of the rejected hypotheses.

Despite their more favorable interpretation, however, the local FDR estimates produced by AdaPT rely on much stronger assumptions than the global FDR control guarantee — namely, that the two-groups model, as well as our specifications for $\pi_1(x)$ and $f_1(p | x)$, must be correct. Instead of using the parametric estimates $\widehat{\text{fdr}}(p_i | x_i)$, we could estimate the local FDR in a moving window of w steps of the AdaPT algorithm:

$$\widehat{\text{fdr}}_{t,w} = \frac{A_t - A_{t+w}}{1 \vee (R_t - R_{t+w})}, \quad \text{or } \widehat{\text{fdr}}_{t,w}^+ = \frac{1 + A_t - A_{t+w}}{1 \vee (R_t - R_{t+w})}.$$

Note that if we take an infinitely large window, we obtain $\widehat{\text{fdr}}_{t,\infty}^+ = \widehat{\text{FDP}}_t$; thus, these estimators adaptively estimate the false discovery proportion for p -values revealed in the next w steps of the algorithm, in much the same way that $\widehat{\text{FDP}}_t$ estimates the false discovery proportion for *all* remaining p -values. It would be interesting to investigate, in future work, what error-control guarantees we might be able to derive by using these estimators.

Acknowledgments

The authors thank Jim Pitman, Ruth Heller and Stefan Wager for helpful discussions.

References

- Rina Foygel Barber and Emmanuel J Candès. Controlling the false discovery rate via knockoffs. *The Annals of Statistics*, 43(5):2055–2085, 2015.
- Rina Foygel Barber and Emmanuel J Candès. A knockoff filter for high-dimensional selective inference. *arXiv preprint arXiv:1602.03574*, 2016.
- Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 289–300, 1995.
- Yoav Benjamini and Yosef Hochberg. Multiple hypotheses testing with weights. *Scandinavian Journal of Statistics*, 24(3):407–418, 1997.
- Richard Berk, Lawrence Brown, Andreas Buja, Kai Zhang, and Linda Zhao. Valid post-selection inference. *The Annals of Statistics*, 41(2):802–837, 2013.
- Sean Davis and Paul S Meltzer. GEOquery: a bridge between the gene expression omnibus (geo) and bioconductor. *Bioinformatics*, 23(14):1846–1847, 2007.
- Edgar Dobriban, Kristen Fortney, Stuart K Kim, and Art B Owen. Optimal multiple testing under a gaussian prior on the effect sizes. *Biometrika*, 102(4):753–766, 2015.
- Lilun Du, Chunming Zhang, et al. Single-index modulated multiple testing. *The Annals of Statistics*, 42(4):1262–1311, 2014.
- Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Leon Roth. Preserving statistical validity in adaptive data analysis. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 117–126. ACM, 2015.
- Bradley Efron. Size, power and false discovery rates. *The Annals of Statistics*, pages 1351–1377, 2007.

- Bradley Efron, Robert Tibshirani, John D Storey, and Virginia Tusher. Empirical bayes analysis of a microarray experiment. *Journal of the American statistical association*, 96 (456):1151–1160, 2001.
- William Fithian, Dennis Sun, and Jonathan Taylor. Optimal inference after model selection. *arXiv preprint arXiv:1410.2597*, 2014.
- Christopher R Genovese, Kathryn Roeder, and Larry Wasserman. False discovery control with p-value weighting. *Biometrika*, 93(3):509–524, 2006.
- Max Grazier G’Sell, Stefan Wager, Alexandra Chouldechova, and Robert Tibshirani. Sequential selection procedures and false discovery rate control. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(2):423–444, 2016.
- James X Hu, Hongyu Zhao, and Harrison H Zhou. False discovery rate control with groups. *Journal of the American Statistical Association*, 2012.
- Nikolaos Ignatiadis, Bernd Klaus, Judith B Zaugg, and Wolfgang Huber. Data-driven hypothesis weighting increases detection power in genome-scale multiple testing. *Nature methods*, 2016.
- Jason D Lee, Dennis L Sun, Yuekai Sun, and Jonathan E Taylor. Exact post-selection inference, with application to the lasso. *The Annals of Statistics*, 44(3):907–927, 2016.
- Lihua Lei and William Fithian. Power of ordered hypothesis testing. In *ICML*, 2016.
- Ang Li and Rina Foygel Barber. Accumulation tests for FDR control in ordered hypothesis testing. *Journal of the American Statistical Association*, (just-accepted):1–38, 2016a.
- Ang Li and Rina Foygel Barber. Multiple testing with the structure adaptive benjamini-hochberg algorithm. *arXiv preprint arXiv:1606.07926*, 2016b.
- John D Storey. A direct approach to false discovery rates. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):479–498, 2002.
- John D Storey. The optimal discovery procedure: a new approach to simultaneous significance testing. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(3):347–368, 2007.
- John D Storey, Jonathan E Taylor, and David Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66 (1):187–205, 2004.
- Wenguang Sun, Brian J Reich, T Tony Cai, Michele Guindani, and Armin Schwartzman. False discovery control in large-scale spatial multiple testing. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(1):59–83, 2015.
- Xiaoying Tian and Jonathan E Taylor. Selective inference with a randomized response. *arXiv preprint arXiv:1507.06739*, 2015.
- John Wilder Tukey. *The collected works of John W. Tukey: Multiple comparisons, 1948-1983*, volume 8. Chapman & Hall/CRC, 1994.
- Daniel Yekutieli. Adjusted bayesian inference for selected parameters. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(3):515–541, 2012.

A EM algorithm details

The fully-observed log-likelihood for the model (10) is

$$\ell(\theta, \beta; y, H, x) = \sum_{i=1}^n H_i \theta' \phi(x_i) - \log \left(1 + e^{-\theta' \phi(x_i)} \right) \quad (14)$$

$$+ \sum_{i=1}^n (-y_i H_i \beta' \phi(x_i) - y_i (1 - H_i)) + \log(\beta' \phi(x_i)) \quad (15)$$

Because some values of y_i and all values of H_i are unknown, we can use the EM algorithm to maximize the partially observed log-likelihood. To simplify estimation, we will proceed as though A_t and R_t are missing, so that the (y_i, H_i) pairs are mutually independent given the predictors. That is, at step t of the AdaPT procedure we attempt to maximize the likelihood of the data $D_t = (x_i, \tilde{p}_{t,i})_{i=1,\dots,n}$ and treating s_t as fixed.

There are four possible values of (b_i, H_i) , with each pair conditionally independent given D_t , and whose probabilities can be efficiently computed for any values of θ and β . Let $r = 0, 1, \dots$ index stages of the EM algorithm (recall t is fixed for the duration of the EM algorithm). For the E-step we compute the expectation of the log-likelihood,

$$\mathbb{E}_{\hat{\theta}^{(r-1)}, \hat{\beta}^{(r-1)}} [\ell(\theta, \beta; y, H, x) | D_t],$$

which amounts to computing the following quantities:

$$\hat{H}_i^{(r)} = \mathbb{E}_{\hat{\theta}^{(r-1)}, \hat{\beta}^{(r-1)}} [H_i | D_t], \quad \text{and} \quad (16)$$

$$\hat{y}_i^{(r,1)} = \mathbb{E}_{\hat{\theta}^{(r-1)}, \hat{\beta}^{(r-1)}} [y_i | D_t, H_i = 1]. \quad (17)$$

where $\hat{\theta}^{(r)}$ and $\hat{\beta}^{(r)}$ denote the current coefficient estimates. For the M-step, we set

$$\hat{\theta}^{(r)}, \hat{\beta}^{(r)} = \arg \max_{\beta, \theta} \mathbb{E}_{\hat{\theta}^{(r-1)}, \hat{\beta}^{(r-1)}} [\ell(\theta, \beta; y, H, x) | D_t] \quad (18)$$

$$= \arg \max_{\beta, \theta} \sum_{i=1}^n \hat{H}_i^{(r)} \theta' \phi(x_i) - \log \left(1 + e^{-\theta' \phi(x_i)} \right) \quad (19)$$

$$+ \sum_{i=1}^n \hat{H}_i^{(r)} \cdot \left(-\hat{y}_i^{(r,1)} \beta' \phi(x_i) + \log(\beta' \phi(x_i)) \right), \quad (20)$$

The optimization above splits into two separate optimization problems, a logistic regression with predictors $\phi(x_i)$ and fractional responses $\hat{H}_i^{(r)}$, and a gamma GLM with predictors $\phi(x_i)$, responses $\hat{y}_i^{(r,1)}$, and weights $\hat{H}_i^{(r)}$. Each of these GLM problems can be solved efficiently using the `glm` function in R. For $r = 0$, we can initialize $\hat{\theta}^{(0)}$ and $\hat{\beta}^{(0)}$ by a simple method with details discussed in Appendix A.2.. Algorithm 2 formalizes the EM algorithm using R pseudocode.

A.1 Derivation of E-step

To fill in the details of the algorithm, we are left to calculate the imputed values $\hat{H}_i^{(r)}$ and $\hat{y}_i^{(r,1)}$ given the parameters θ and β . Denote π_i and μ_i by

$$\pi_i = \left(1 + e^{-\phi(x)' \theta} \right)^{-1}, \quad \mu_i = (\phi(x)' \beta)^{-1}$$

We distinguish two cases. In the first case, $\tilde{p}_{t,i}$ is a number (instead of a set), then by definition $\tilde{p}_{t,i} = p_i$. Thus,

$$\hat{H}_i^{(r)} = \mathbb{P}(H_i = 1 | p_i = \tilde{p}_{t,i}) = \frac{\pi_i \cdot \frac{1}{\mu_i} \tilde{p}_{t,i}^{\frac{1}{\mu_i} - 1}}{\pi_i \cdot \frac{1}{\mu_i} \tilde{p}_{t,i}^{\frac{1}{\mu_i} - 1} + 1 - \pi_i}, \quad (21)$$

Algorithm 2 EM algorithm to estimate $\pi_1(\cdot)$ and $\mu(\cdot)$ based on $D_t = (x_i, \tilde{p}_{t,i})_{i=1,\dots,n}$

Input: data D_t , number of iterations m , initialization $\hat{\theta}^{(0)}, \hat{\beta}^{(0)}$;

for $r = 1, 2, \dots, m$ **do**

 (*E-step*):

$$\hat{H}_i^{(r)} \leftarrow \mathbb{E}_{\hat{\theta}^{(r-1)}, \hat{\beta}^{(r-1)}}[H_i \mid D_t], \quad i = 1, \dots, n;$$

$$\hat{y}_i^{(r,1)} \leftarrow \mathbb{E}_{\hat{\theta}^{(r-1)}, \hat{\beta}^{(r-1)}}[y_i \mid D_t, H_i = 1], \quad i = 1, \dots, n;$$

 (*M-step*):

$$\hat{\theta}^{(r)} \leftarrow \text{glm}(\hat{H}^{(r)} \sim \phi(x), \text{family} = \text{binomial});$$

$$\hat{\beta}^{(r)} \leftarrow \text{glm}(\hat{y}^{(r,1)} \sim \phi(x), \text{family} = \text{gamma}, \text{weights} = \hat{H}^{(r)});$$

end for

Output: $\hat{\pi}_1(x) = \left(1 + e^{-\phi(x)' \hat{\theta}^{(m)}}\right)^{-1}$, $\hat{\mu}(x) = \left(\phi(x)' \hat{\beta}^{(m)}\right)^{-1}$.

and

$$\hat{y}_i^{(r,1)} = \mathbb{E}(y_i H_i \mid \tilde{p}_{t,i}) / \hat{H}_i^{(r)} = y_i. \quad (22)$$

In the second case, $\tilde{p}_{t,i}$ is a two-elements set. Denote the smaller element by $q_{t,i}$, i.e. $q_{t,i} = \min\{p_i, 1 - p_i\}$, then conditioning on $\tilde{p}_{t,i} = \{q_{t,i}, 1 - q_{t,i}\}$, p_i is supported on $\{q_{t,i}, 1 - q_{t,i}\}$ with

$$\mathbb{P}(p_i = q_{t,i} \mid \tilde{p}_{t,i}, H_i = 1) = \frac{q_{t,i}^{\frac{1}{\mu_i} - 1}}{q_{t,i}^{\frac{1}{\mu_i} - 1} + (1 - q_{t,i})^{\frac{1}{\mu_i} - 1}} \quad (23)$$

and

$$\mathbb{P}(p_i = 1 - q_{t,i} \mid \tilde{p}_{t,i}, H_i = 1) = \frac{(1 - q_{t,i})^{\frac{1}{\mu_i} - 1}}{q_{t,i}^{\frac{1}{\mu_i} - 1} + (1 - q_{t,i})^{\frac{1}{\mu_i} - 1}}. \quad (24)$$

By Bayes' formula, we can also derive the conditional distribution of H_i given $\tilde{p}_{t,i}$:

$$\hat{H}_i^{(r)} = \mathbb{P}(H_i = 1 \mid \tilde{p}_{t,i}) = \frac{\pi_i \cdot \frac{1}{\mu_i} \left(q_{t,i}^{\frac{1}{\mu_i} - 1} + (1 - q_{t,i})^{\frac{1}{\mu_i} - 1} \right)}{\pi_i \cdot \frac{1}{\mu_i} \left(q_{t,i}^{\frac{1}{\mu_i} - 1} + (1 - q_{t,i})^{\frac{1}{\mu_i} - 1} \right) + 2(1 - \pi_i)}. \quad (25)$$

As a consequence of (23) - (25),

$$\begin{aligned} \hat{y}_i^{(r,1)} &= \frac{1}{\hat{H}_i^{(r)}} \cdot (\log q_{t,i} \cdot \mathbb{P}(H_i = 1, p_i = q_{t,i} \mid \tilde{p}_{t,i}) + \log(1 - q_{t,i}) \cdot \mathbb{P}(H_i = 1, p_i = 1 - q_{t,i} \mid \tilde{p}_{t,i})) \\ &= \frac{1}{\hat{H}_i^{(r)}} \cdot \frac{\pi_i \cdot \frac{1}{\mu_i} \left(q_{t,i}^{\frac{1}{\mu_i} - 1} \log q_{t,i} + (1 - q_{t,i})^{\frac{1}{\mu_i} - 1} \log(1 - q_{t,i}) \right)}{\pi_i \cdot \frac{1}{\mu_i} \left(q_{t,i}^{\frac{1}{\mu_i} - 1} + (1 - q_{t,i})^{\frac{1}{\mu_i} - 1} \right) + 2(1 - \pi_i)} \\ &= \frac{q_{t,i}^{\frac{1}{\mu_i} - 1} \log q_{t,i} + (1 - q_{t,i})^{\frac{1}{\mu_i} - 1} \log(1 - q_{t,i})}{q_{t,i}^{\frac{1}{\mu_i} - 1} + (1 - q_{t,i})^{\frac{1}{\mu_i} - 1} + 2(1 - \pi_i)}. \end{aligned} \quad (26)$$

A.2 Initialization

Another important issue is the initialization. The formulae of $\hat{H}_i^{(r)}$ and $\hat{y}_i^{(r,1)}$ requires estimates of $\pi(x_i)$ and $\mu(x_i)$. In step 0 when no information can be obtained, we propose

a simple method by imputing random guess as follows. First we obtain an initial guess of $\pi_1(x_i)$. Let $J_i = I(\tilde{p}_{t,i} \text{ contains two elements})$, then we observe that

$$\begin{aligned} \mathbb{E}J_i &= \mathbb{P}(p_i \notin [s_0(x_i), 1 - s_0(x_i)]) \geq (1 - \pi_1(x_i))(1 - 2s_0(x_i)) \\ \implies \pi_1(x_i) &\geq \mathbb{E}\left(1 - \frac{J_i}{1 - 2s_0(x_i)}\right). \end{aligned} \quad (27)$$

Let

$$\tilde{J}_i = 1 - \frac{J_i}{1 - 2s_0(x_i)} \quad (28)$$

and then we fit a linear regression on \tilde{J}_i with covariates $\phi(x_i)$, denoted by $\hat{\pi}_0(x_i)$ and then truncate $\hat{\pi}_0(x_i)$ at 0 and 1 to obtain an initial guess of $\pi_1(x_i)$, i.e.

$$\pi_0(x_i) = (\hat{\pi}_0(x_i) \vee 0) \wedge 1. \quad (29)$$

(28) implies that $\hat{\pi}_0(\cdot)$ is a conservative estimate of $\pi_1(\cdot)$. This is preferred to an anti-conservative estimate since the latter might cause over-fitting.

Then we obtain an initial guess of $\mu(x_i)$ based on $\hat{\pi}_0(\cdot)$ by imputing p'_i . If $p_i \in [s_0(x_i), 1 - s_0(x_i)]$, then $\tilde{p}_{t,i} = p_i$ and hence we can use it directly. Otherwise, we only know that $p_i \in \tilde{p}_{t,i} = \{q_{t,i}, 1 - q_{t,i}\}$. If p_i is null, then it should be uniform on $\{q_{t,i}, 1 - q_{t,i}\}$; if p_i is non-null, then it should more likely to be $q_{t,i}$ since $q_{t,i} < 1 - q_{t,i}$. Thus, we impute p'_i as

$$p'_i = \begin{cases} q_{t,i} & \text{with probability } \hat{\pi}_0 + \frac{1}{2}(1 - \hat{\pi}_0) \\ 1 - q_{t,i} & \text{with probability } \frac{1}{2}(1 - \hat{\pi}_0) \end{cases} \quad (30)$$

Then we can fit an unweighted Gamma GLM on $-\log \tilde{p}_i$ with covariates $\phi(x_i)$ and inverse link to obtain an initial guess of $\mu(x_i)$, denoted by $\mu_0(x_i)$.

B Technical Proofs

Proof of Theorem 2. Assume $f_0(p | x_i) \leq B$. Denote $\eta_i = \nu(\{x_i\})$ and $s = (s_1, \dots, s_n) \triangleq (s(x_1), \dots, s(x_n))$, then the objective function of (7)

$$\int_{\mathcal{X}} -F_1(s(x)|x)\pi_1(x)\nu(dx) = -\sum_{i=1}^n \eta_i F_1(s_i|x_i)\pi_1(x_i)$$

is a convex function of s by condition (i) and the constraint function

$$\begin{aligned} g(s) &\triangleq \int_{\mathcal{X}} \left\{ -\alpha F_1(s(x)|x)\pi_1(x) + (1 - \alpha)F_0(s(x)|x)(1 - \pi_1(x)) \right\} \nu(dx) \\ &= \sum_{i=1}^n \eta_i (-\alpha F_1(s_i|x_i)\pi_1(x_i) + (1 - \alpha)F_0(s_i|x_i)(1 - \pi_1(x_i))) \end{aligned}$$

is also a convex function of s by condition (i). To establish the necessity of KKT condition, it is left to prove the Slater's condition, i.e. there exists a \bar{s} , such that for any $s \in B(\bar{s}, \delta)$ for some $\delta > 0$, the constraint inequality holds, i.e. $g(s) \leq 0$, and $g(\bar{s}) < 0$. By condition (ii), WLOG we assume $\widehat{\text{fdr}}(0|x_1) < \alpha$ with $\nu(\{x_1\}) = \eta_1 > 0$. Let $\bar{s} \in \mathbb{R}^n$ with

$$\bar{s}_1 = 2\epsilon, \bar{s}_2 = \dots = \bar{s}_n = \epsilon \cdot (2B)^{-1} \eta_1 \Delta$$

where

$$\Delta = f(0|x_1) \left(\alpha(1 - \widehat{\text{fdr}}(0|x_1)) - (1 - \alpha)\widehat{\text{fdr}}(0|x_1) \right) > 0.$$

Let $\delta = \min\{\bar{s}_2, \epsilon\}$. We will show that for any $s \in B(\bar{s}, \delta)$, $g(s) < 0$. In fact, for any $s \in B(\bar{s}, \delta)$, we have

$$s_1 \in [\epsilon, 3\epsilon], \quad s_i \leq \epsilon \cdot B^{-1} \eta_1 \Delta.$$

By mean-value theorem, there exists $\tilde{s}_1, \dots, \tilde{s}_n \in [0, \epsilon]$, such that

$$\begin{aligned}
g(s) &= s_1 \eta_1 (-\alpha f_1(\tilde{s}_1|x_1)\pi_1(x_1) + (1-\alpha)f_0(\tilde{s}_1|x_1)(1-\pi_1(x_1))) \\
&\quad + \sum_{i=2}^n s_i \eta_i (-\alpha f_1(\tilde{s}_i|x_i)\pi_1(x_i) + (1-\alpha)f_0(\tilde{s}_i|x_i)(1-\pi_1(x_i))) \\
&= s_1 \eta_1 (-\alpha f_1(0|x_1)\pi_1(x_1) + (1-\alpha)f_0(0|x_1)(1-\pi_1(x_1)) + o(1)) \\
&\quad + \sum_{i=2}^n s_i \eta_i (-\alpha f_1(\tilde{s}_i|x_i)\pi_1(x_i) + (1-\alpha)f_0(\tilde{s}_i|x_i)(1-\pi_1(x_i))) \\
&= -s_1 \eta_1 \Delta + o(s_1 \eta_1) + \sum_{i=2}^n s_i \eta_i (-\alpha f_1(\tilde{s}_i|x_i)\pi_1(x_i) + (1-\alpha)f_0(\tilde{s}_i|x_i)(1-\pi_1(x_i))) \\
&\leq -\epsilon \eta_1 \Delta + o(\epsilon) + \epsilon \eta_1 \Delta \sum_{i=2}^n \eta_i \\
&\leq -\epsilon \eta_1^2 \Delta + o(\epsilon).
\end{aligned}$$

Thus, for sufficiently small ϵ , $g(s) < 0$ for all $s \in B(\bar{s}, \delta)$ and hence the Slater's condition is satisfied. \square

Proof of Lemma 1. We assume $\rho < 1$ (otherwise the result is trivial). Following Barber and Candès (2016), we introduce the random set $\mathcal{A} \subseteq \{1, \dots, n\}$ with

$$\mathbb{P}(i \in \mathcal{A} \mid \mathcal{G}_{-1}) = \frac{1 - \rho_i}{1 - \rho},$$

conditionally independent for $i = 1, \dots, n$, and construct conditionally i.i.d. Bernoulli variables q_1, \dots, q_n with $\mathbb{P}(q_i = 1 \mid \mathcal{G}_{-1}) = \rho$. Then we can define

$$b_i = q_i \mathbf{1}\{i \in \mathcal{A}\} + \mathbf{1}\{i \notin \mathcal{A}\},$$

which by construction gives $\mathbb{P}(b_i = 1 \mid \mathcal{G}_{-1}) = \rho_i$ almost surely, with the b_i conditionally independent.

To ensure that \mathcal{C}_t decreases by at most a single element in each step, we introduce intermediate steps: for integers $t \geq 0$, $1 \leq i \leq n$ define

$$\mathcal{C}_{t+i/n} = \mathcal{C}_t \cup \{j \leq i : j \in \mathcal{C}_{t+1}\}.$$

Next, define the augmented filtration

$$\mathcal{G}_t^A = \sigma \left(\mathcal{G}_{-1}, \mathcal{A}, \mathcal{C}_t, (b_i)_{i \notin \mathcal{C}_t \cap \mathcal{A}}, \sum_{i \in \mathcal{C}_t \cap \mathcal{A}} b_i \right) \supseteq \mathcal{G}_t,$$

for both integer and fractional values of t . Note $\mathcal{C}_{t+1/n}$ is measurable with respect to \mathcal{C}_t .

Next, define

$$U_t = \sum_{i \in \mathcal{C}_t \cap \mathcal{A}} b_i, \quad V_t = \sum_{i \in \mathcal{C}_t \cap \mathcal{A}} 1 - b_i, \quad \text{and} \quad Z_t = \frac{1 + |\mathcal{C}_t|}{1 + U_t}.$$

Finally, we observe that $(b_i)_{i \in \mathcal{C}_t \cap \mathcal{A}} = (q_i)_{i \in \mathcal{C}_t \cap \mathcal{A}}$ are exchangeable with respect to \mathcal{G}_t^A , with the random vector distributed uniformly over configurations summing to U_t .

There are three cases: First, if $\mathcal{C}_{t+1/n} = \mathcal{C}_t$ then $\mathbb{E}[Z_{t+1/n} \mid \mathcal{G}_t^A] = Z_t$. Otherwise, $(U_{t+1/n}, V_{t+1/n})$ is either $(U_t - 1, V_t)$ or $(U_t, V_t - 1)$ with conditional probabilities proportional to U_t and V_t . Second, if $U_t = 0$ then $Z_{t+1/n} = V_t \leq Z_{t+1/n} = 1 + V_t$. Third, if $U_t > 0$, then

$$\begin{aligned}
\mathbb{E}[Z_{t+1/n} \mid \mathcal{G}_t^A] &= \frac{U_t + V_t}{1 + U_t} \cdot \frac{V_t}{U_t + V_t} + \frac{U_t + V_t}{U_t} \cdot \frac{U_t}{U_t + V_t} \\
&= \frac{V_t}{1 + U_t} + 1 = Z_t.
\end{aligned}$$

In all three cases, the conditional expectation of $Z_{t+1/n}$ is smaller than Z_t ; thus, Z_t is a supermartingale with respect to the filtration \mathcal{G}_t^A . Because \hat{t} is also a stopping time with respect to the filtration $(\mathcal{G}_t^A)_{t=0,1/n,2/n,\dots}$ (but one which can only take integer values), we have

$$\mathbb{E} \left[\frac{1 + |\mathcal{C}_{\hat{t}}|}{1 + \sum_{i \in \mathcal{C}_{\hat{t}}} b_i} \middle| \mathcal{G}_{-1} \right] \leq \mathbb{E} [Z_{\hat{t}} | \mathcal{G}_{-1}] \quad (31)$$

$$\leq \mathbb{E} [Z_0 | \mathcal{G}_{-1}] \quad (32)$$

$$\leq \rho^{-1}. \quad (33)$$

See Barber and Candès (2015) for justification of the last inequality. Marginalizing over \mathcal{G}_{-1} we obtain the result. \square

C Sensitivity Analysis

Our implementation of AdaPT for AdaPT involves two parameters: the number of knots N and the tolerance parameter δ . It is conducive to examine the robustness of AdaPT on the parameters. Figure 6 and Figure 7 shows the number of rejections with different choice of N and δ . When $\alpha < 0.2$, it is clear that AdaPT is fairly robust to the choice of number of knots. For large α , different number of knots leads to indifferent performance with moderately informative ordering. With original ordering, choosing a small number of knots slightly degrades the performance while with highly informative ordering, choosing a large number of knots does so. To explain the phenomena, note that for the former, the signals (non-nulls) are scattered and form small clusters as shown in Figure 3. Thus, using more knots is better for capturing the local information and identifying the clusters. By contrast, the latter case has an informative ordering and the signals concentrate on the beginning of the list as shown in Figure 5. For this reason, a few knots is sufficient for revealing the signals while a large number of knots might cause overfitting due to the local oscillation around the knots. In summary, if the target FDR level is not too large as in many applications, AdaPT is insensitive to the number of knots. For large target FDR level, we suggest increasing the number of knots when the signals are dispersed while decreasing the number of knots when the signals are concentrated. A practical recommendation is setting a knot per 1000-2000 samples.




Figure 6: Number of discoveries by AdaPT with different number of knots at a range of target FDR levels α from 0.01 to 0.30. Each panel plots the results for an ordering, ranging from no ordering to highly informative.

As for δ , the top panels of Figure 7 show that a small δ generally leads to more rejections, which is heuristically correct since the procedure is more patient. However, the number of rejections is insensitive to δ when α is small ($\alpha < 0.2$) and δ is not too large ($\delta < 0.1$). In




Figure 7: Number of discoveries (top) and number of iterations (bottom) by AdaPT with different patience parameter δ at a range of target FDR levels α from 0.01 to 0.30. Each panel plots the results for an ordering, ranging from no ordering to highly informative.

addition, the bottom panels of Figure 7 show the number of iterations to reach the level α . It is seen that a small δ significantly increases the computation cost, which is almost inversely proportional to δ . To balance the computation cost and the power, we suggest fixing δ in the range $[0.05, 0.1]$.

Another aspect worth to explore is the choice of the covariate x . In section 5, we choose the order derived from another sequences of p-values, denoted by \tilde{p}_i . A natural alternatives of the covariates include the \tilde{p}_i and $\text{logit}(\tilde{p}_i) = \log \frac{\tilde{p}_i}{1-\tilde{p}_i}$. Figure 8 shows the number of rejections of AdaPT with three choices of x and a same set of parameter N and δ as in section 5. The performance is slightly worse using \tilde{p}_i but slightly better using $\text{logit}(\tilde{p}_i)$ than that using the order. However, the difference is too small to witness for most FDR levels and we believe it is not significant.

Gene/drug response data: results (with different x)

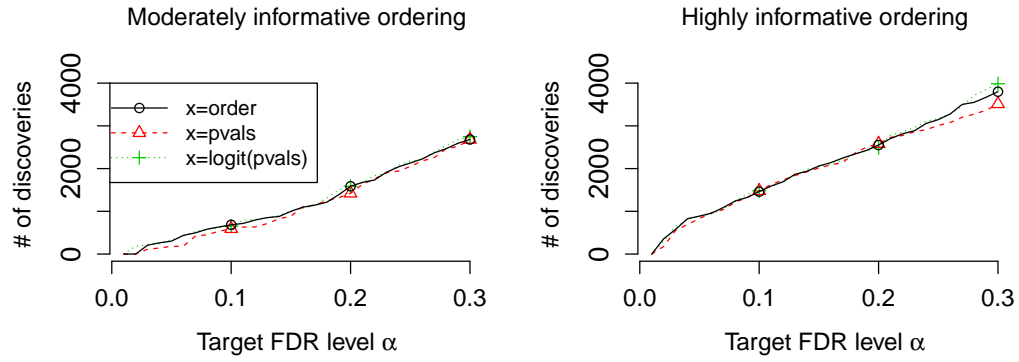


Figure 8: Number of discoveries by AdaPT with choices of covariates x at a range of target FDR levels α from 0.01 to 0.30. Each panel plots the results for an ordering, ranging from no ordering to highly informative.